# Montecito based Systems

**Daniel Reich**
**BCS EMEA, HP**


Montecito

# Intel® Server Platform Roadmap

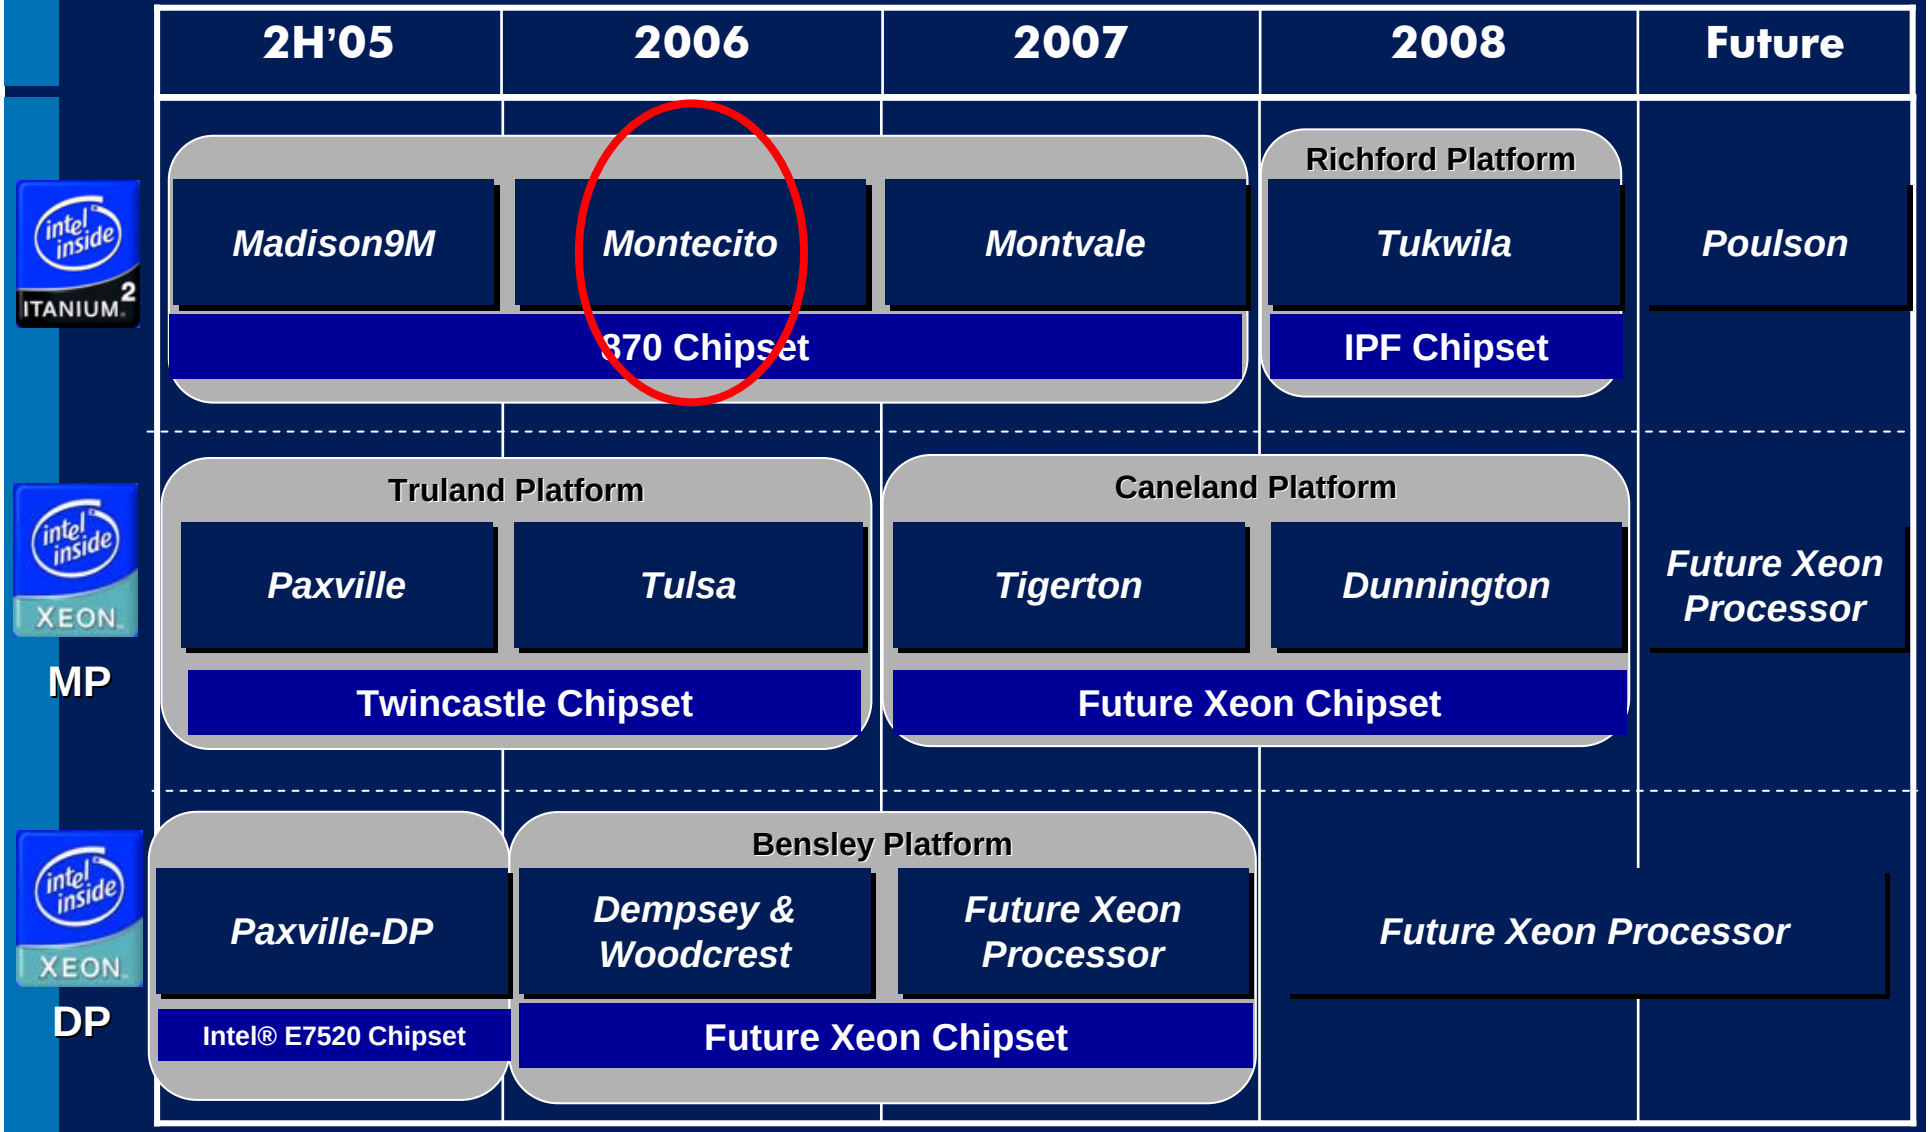




10 November 2006

# Introducing Montecito

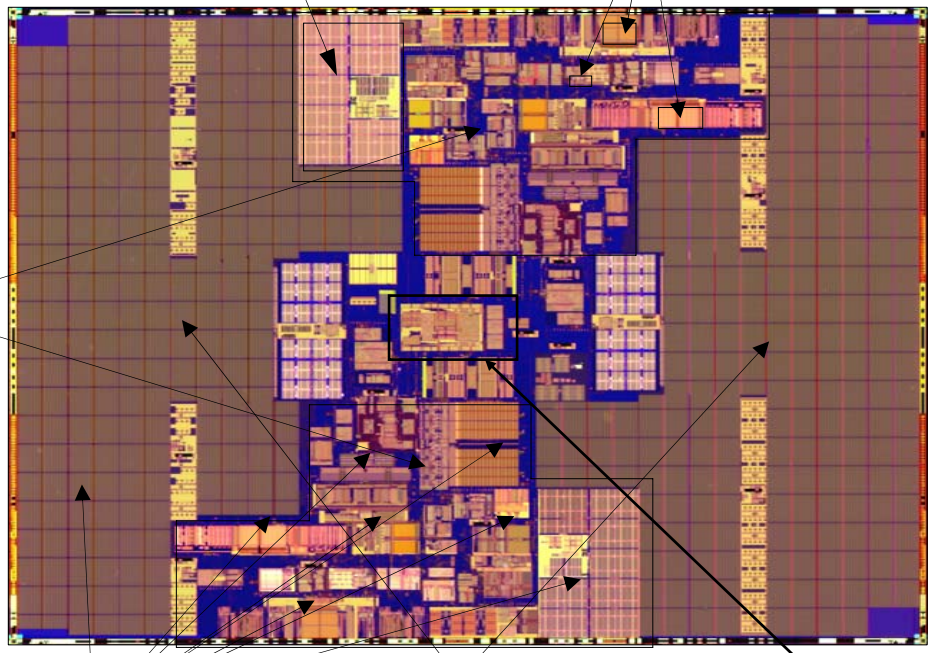

# Quick Montecito Description

- Terminology tutorial:
  - Processor (was socket),
  - Core
  - Hyperthread (or hardware thread), LCPU in HP-UX
  - 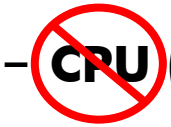

(it's too ambiguous)
- 90 nm version of Madison
- 2 cores per processor
- 1-2 hyperthreads per core
  - Up to 4 hyperthreads (or LCPUs) per processor
- L3 cache: 24MB (12MB per core)
- 130W die, 170W envelope
- E.g. sx2000 with Montecito – 64 processors , 2 cores per processor, up to 2 hyperthreads per core means 256 max hardware threads (or LCPUs)
- 11.23 (without HT)—>11.31 (with HT)—>

# Montecito overview

- Technology
  - 1.6 GHz core speed
  - 130 W
  - 90 nm
- Cache (26.5MB total)
  - Maximum 12 MB L3 cache per core (24MB total, non-shared)
  - Maximum L3 associativity of 12
  - Minimum L3 access time of 14 cycles
  - The L2 cache has been split into separate L2I (1MB) and L2D (256KB) caches per core
  - 16 KB L1I, 16KB L1D per core
- RAS + Manageability features
  - Additional ECC and parity for structures
  - ECC add in L2T tags, and parity added to L1I TLB
- First die with over 1 billion transistors (1.72B total)

# Montecito overview

- Multi-threading
  - 2 cores / socket or cpu; 2 threads / core
  - Processor core and memory hierarchy is based on Itanium® 2
    - With improvements in control, data paths and resources
  - McKinley bus compatible
  - CMT /SOEMT (Switch- On- Event) Multi-threading
    - Switches occur on event (on cache misses, hints, timeout, …)
- Pellston
  - Mapping out (disabling) of cache lines with errors to avoid hard errors that would crash the partition
- VT-i (Silverdale)
  - First IPF processor to implement Virtual Machine hooks to ease the implementation of VM monitors (Vanderpool/VT on x86)

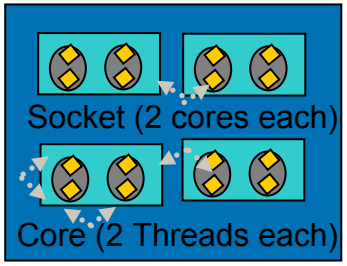


Chipsets

# Chipsets for the new server

- For entry level servers (1-4 Processor) :

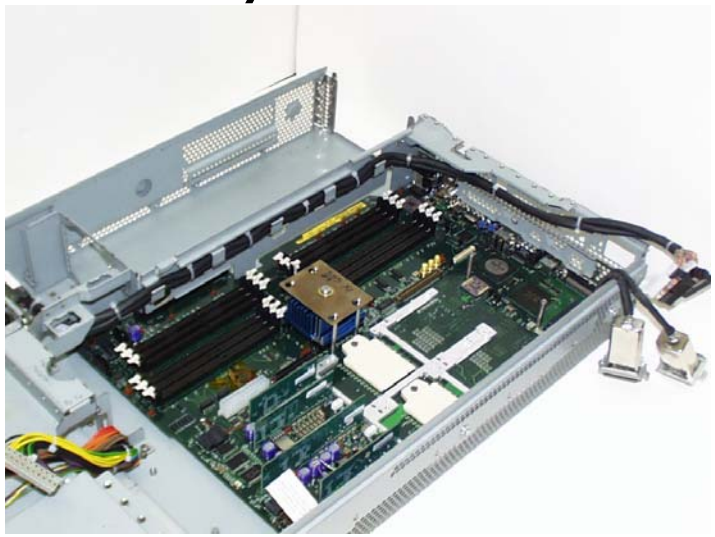

The zx2 chipset

- For mid range and high end servers

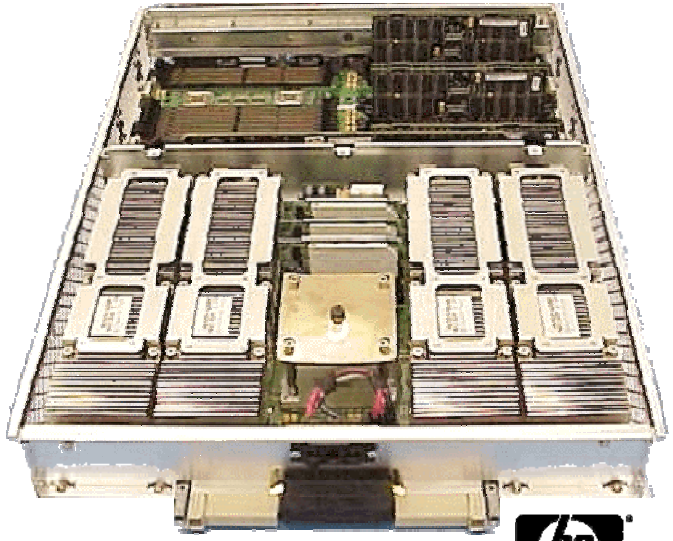

The sx2000 chipset

# zx2 chipset components

The zx2 chipset contains three components

- zx2 memory & I/O controller "Titan"
  - Connects to processor bus
  - Contains memory controller
  - Contains I/O cache controller
- zx2 I/O adapters
  - "Mercury" PCI, PCI-X 133MHz (zx1 I/O adapter)
  - "Gemini" PCI-X DDR 266MHz
  - "Apollo" PCI-Express
- zx2 scalable memory expander "Agena"
  - Optional component used to:
    - (Primarily) Increase memory capacity
    - Increase memory bandwidth


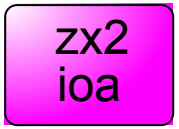


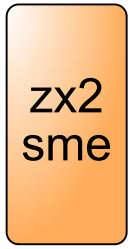

# zx2 : continued performance leadership

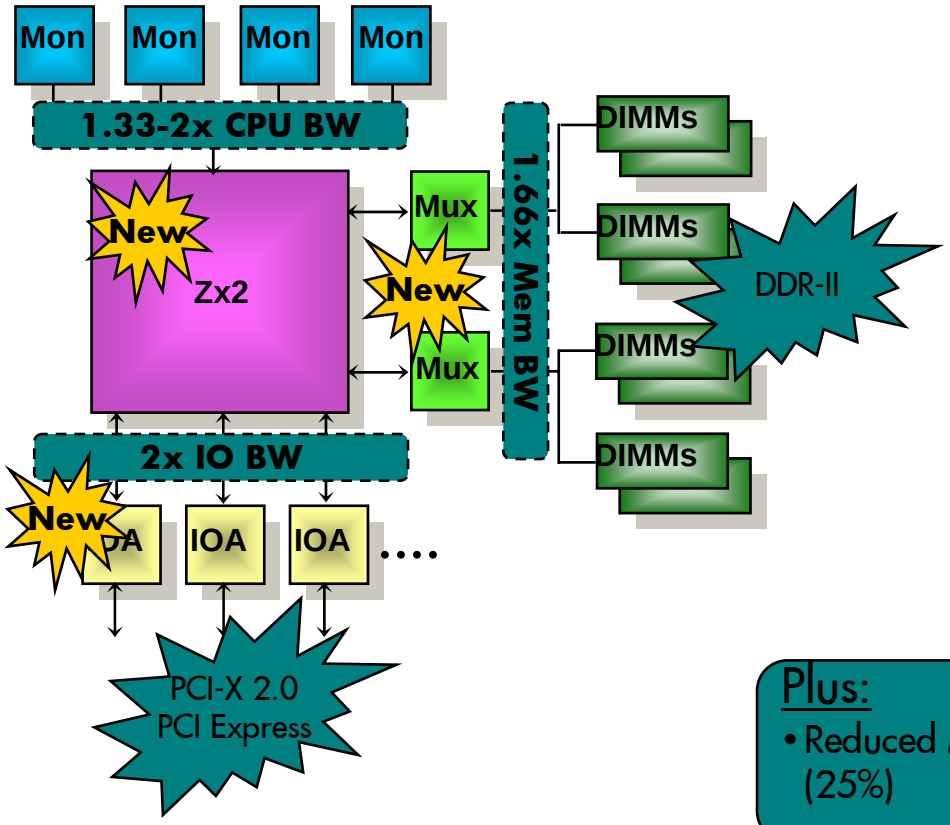


Plus:

- Reduced Memory Latency (25%)



invent

# Comparison: HP zx1 and zx2 chipsets

| Feature | zx1 | zx2 |
|---|---|---|
| FSB BW – 3 load | 6.4-8.5 GB/s | 10.6 GB/s |
| FSB BW – 5 load | 6.4 GB/s | 8.5 GB/s |
| Memory technology | DDR 266 | DDR-II 444/533 |
| Memory BW (total) | Direct Attach: 8.5 GB/s<br>w/expander: 12.8 GB/s | Direct Attach: 14.2 – 17.0 GB/s<br>w/expander: 12.8 – 17.0 GB/s |
| Latency (pin to pin) | Direct Attach: 95ns<br>w/expander: 115ns | Direct Attach: 60ns<br>w/expander: 78ns |
| Max # DIMMs | 16 direct attach<br>64 w/expander | 12 direct attach<br>48 w/expander |
| Memory capacity | 256GB | 768GB (capable) |
| Memory scrubber | Software | Hardware |
| DRAM chip sparing | Single | Double |
| IO technology | PCI, PCI-X, AGP 8X | (add) PCI-X 2.0, PCI-Express |
| IO bandwidth (total) | 4 GB/s | 8-10 GB/s |
| DMA read latency | 700ns | 400ns |

NM1

# sx2000

**Maximizes application performance from 4 to 128 processors**

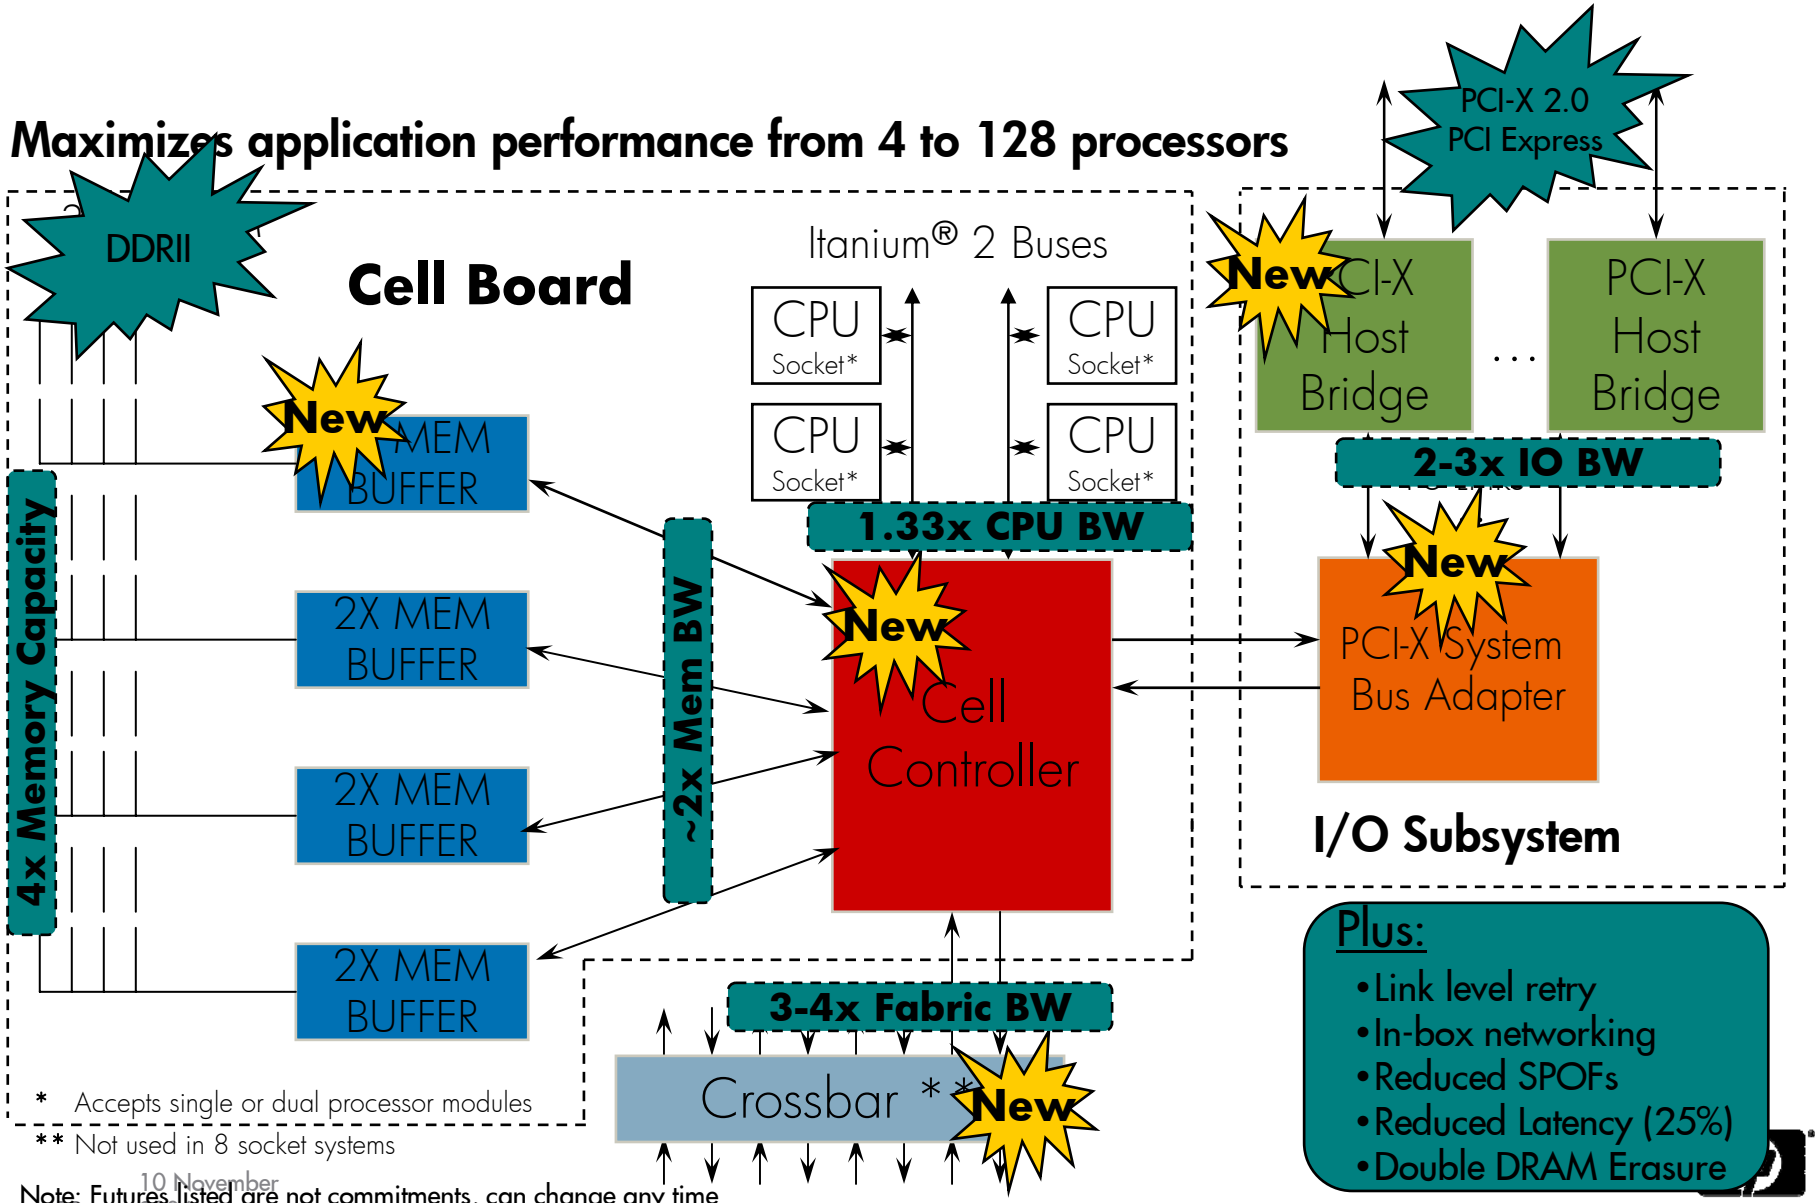

## Slide 13

**NM1**

Focus on a textual

ord on delivered value to the market

--- latency improvments
coherency improvements
RAS/HA...
More redundncy ---

PULL SLIDE FROM GREG'S TAF DECK
Neil MacDonald; 01.02.2005

# What's new in sx2000 systems

- Memory
  - DDR2 memory technology
  - Increased memory address space
  - Enhanced DRAM ECC
- I/O
  - PCI-X Mode 2
  - PCI Express in future release
- System Interconnect
  - New link technology for fabric and I/O
  - Upgraded crossbar topology
- Performance improvements
  - Increased bandwidths, reduced latencies
- Networking over the backplane
- I/O Write coalescing
- Fail-over system clock for high reliability

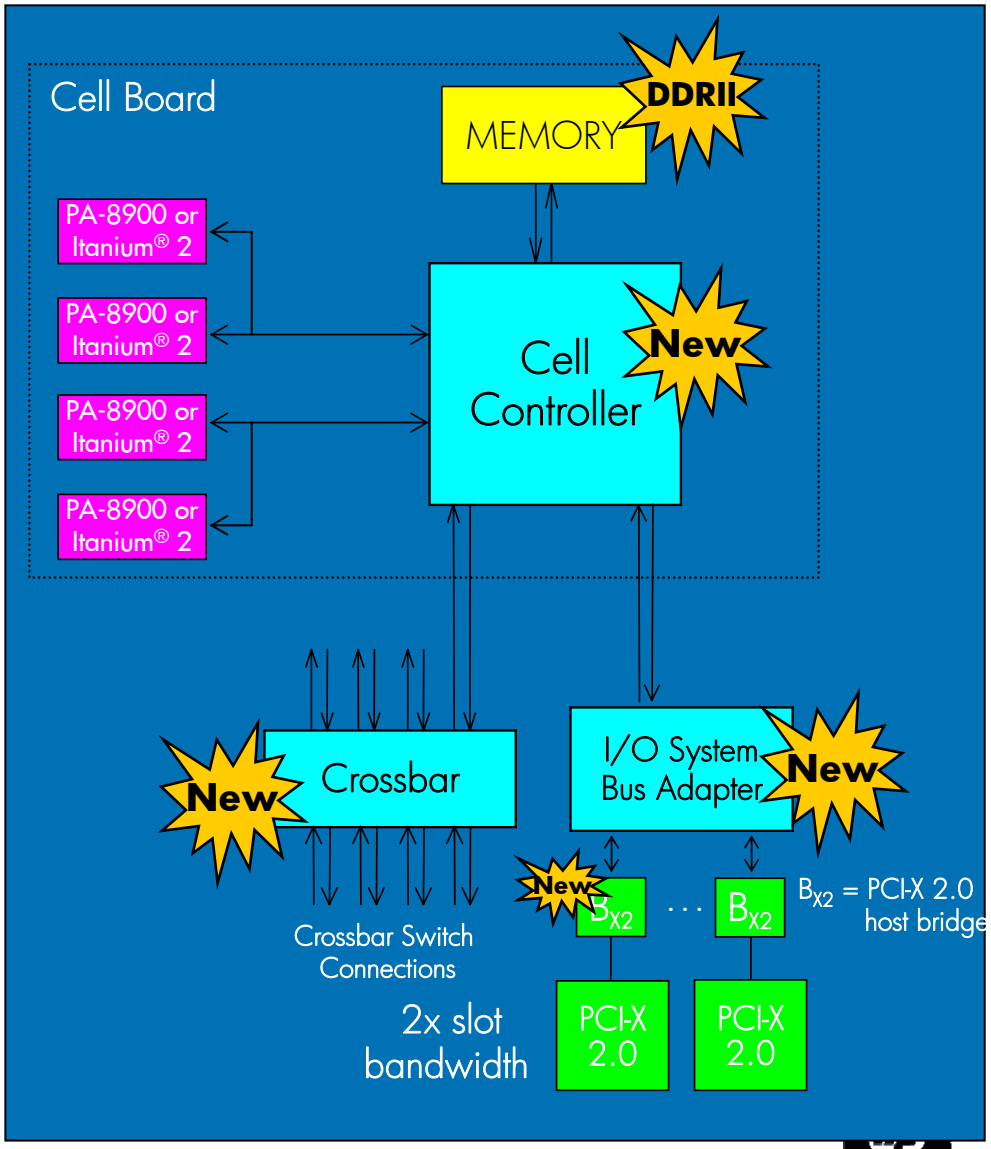

# sx1000 in Superdome

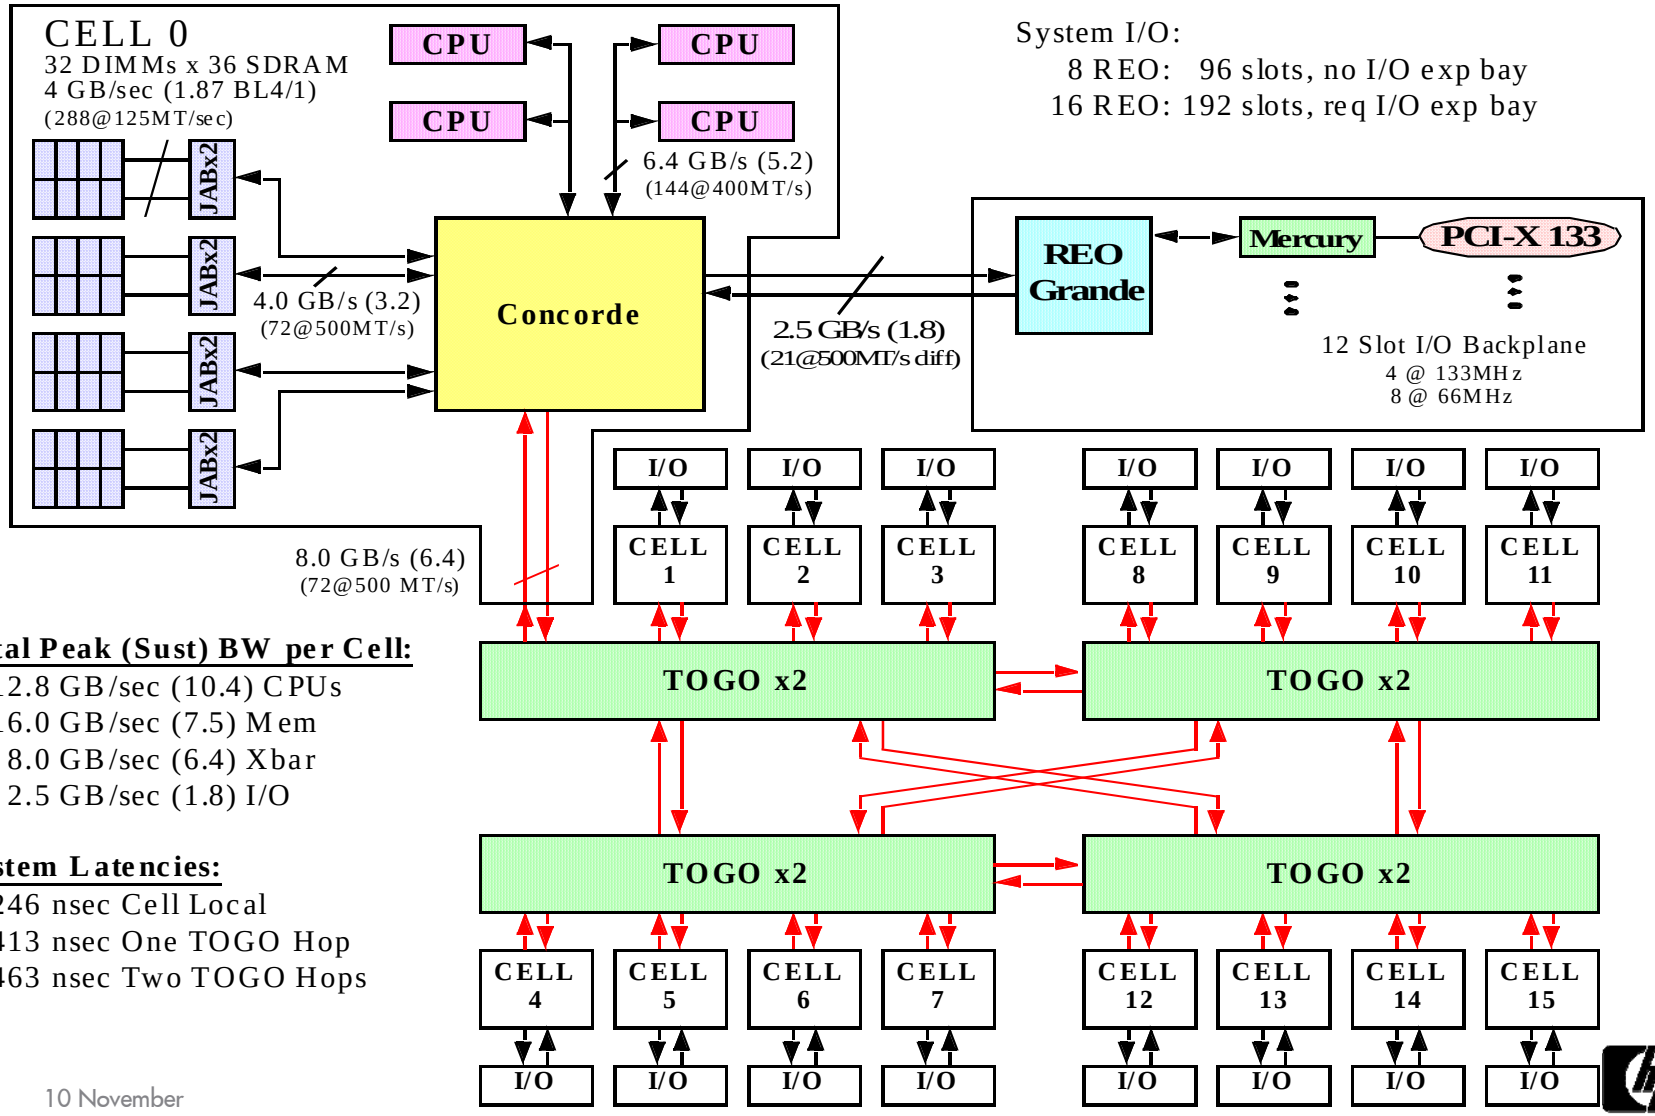


System I/O:
8 REO: 96 slots, no I/O exp bay
16 REO: 192 slots, req I/O exp bay

**Total Peak (Sust) BW per Cell:**
12.8 GB/sec (10.4) CPUs
16.0 GB/sec (7.5) Mem
8.0 GB/sec (6.4) Xbar
2.5 GB/sec (1.8) I/O

**System Latencies:**
246 nsec Cell Local
413 nsec One TOGO Hop
463 nsec Two TOGO Hops

# sx2000 in Superdome

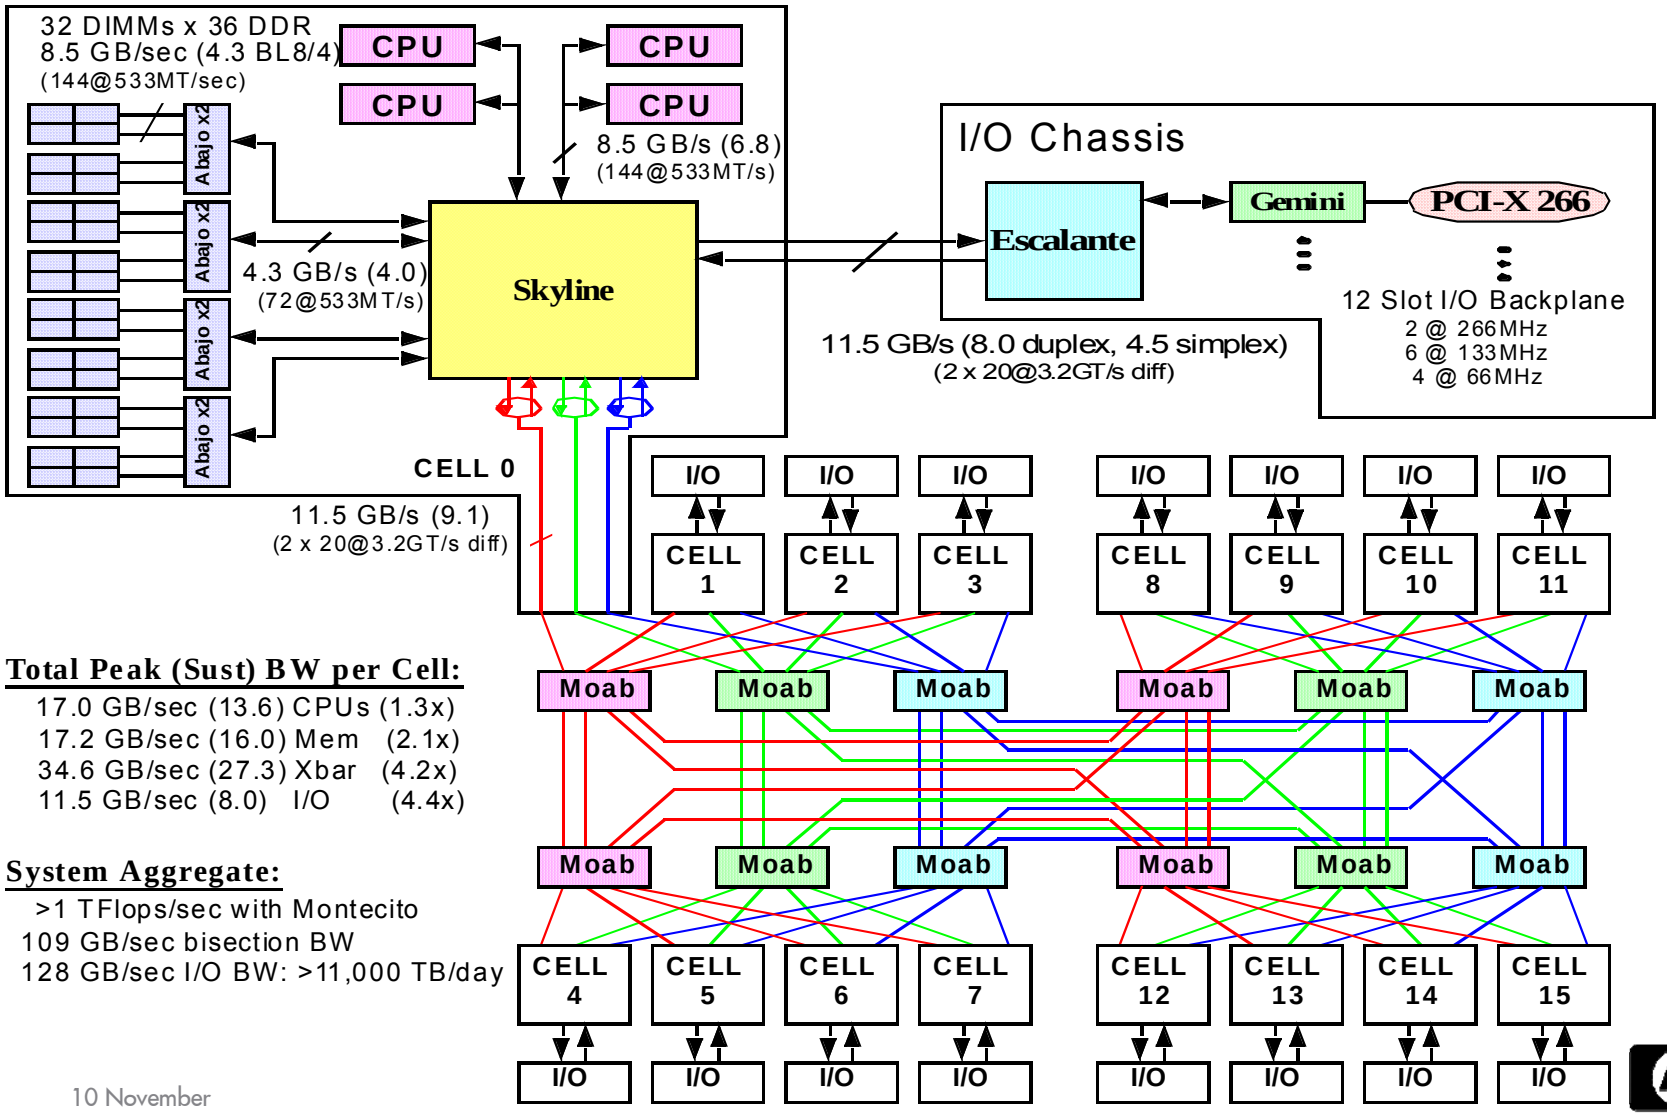


**Total Peak (Sust) BW per Cell:**

17.0 GB/sec (13.6) CPUs (1.3x)
17.2 GB/sec (16.0) Mem (2.1x)
34.6 GB/sec (27.3) Xbar (4.2x)
11.5 GB/sec (8.0) I/O (4.4x)

**System Aggregate:**

>1 TFlops/sec with Montecito
109 GB/sec bisection BW
128 GB/sec I/O BW: >11,000 TB/day


Montecito based Server

# HP Integrity and HP 9000 Server Roadmap

Revision: OB.07.06.0 July-06

| Current offering | 2H/2006 | 2007 | 2008 | Future |
|---|---|---|---|---|
|  HP 9000 Superdome HP-UX 11iv1, v2 | | New sx2000 Chipset DDR-II CPU: PA-8900 OS: HP-UX 11iv1, v2, v3 End of PA-RISC releases | New Server & Chipset CPU: Tukwila OS: HP-UX, Windows Linux, OpenVMS | New CPU: Poulson |
| HP Integrity Superdome HP-UX, Windows, Linux, OpenVMS | CPU: Itanium2 "Montecito" OS: HP-UX 11iv2,v3, Windows, Linux, OpenVMS | CPU: Itanium2 "Montvale" PCI-E OS: HP-UX 11iv2,v3, Windows, Linux, OpenVMS | | |
|  HP 9000 rp7420, rp8420 HP-UX 11iv1, v2 | | New sx2000 Chipset DDR-II CPU: PA-8900 OS: HP-UX 11iv1, v2, v3 End of PA-RISC releases | New 16P Servers & Chipset CPU: Tukwila OS: HP-UX, Windows Linux, OpenVMS | |
|  HP Integrity rx7640, rx8640 HP-UX, Windows, Linux, OpenVMS | CPU: Itanium2 "Montecito" OS: HP-UX 11iv2,v3, Windows, Linux, OpenVMS | CPU: Itanium2 "Montvale" PCI-E OS: HP-UX 11iv2,v3, Windows, Linux, OpenVMS | New 8P Servers & Chipset CPU: Tukwila OS: HP-UX, Windows Linux, OpenVMS | |
|  HP Integrity rx4640 HP-UX, Windows, Linux, OpenVMS | New 4p/8c Server & Chipset CPU: Itanium2 "Montecito" DDR-II PCI-E OS: HP-UX, Windows, Linux, OpenVMS<br>CPU: Itanium2 "Montecito" OS: HP-UX 11iv2, v3*, Windows, Linux, OpenVMS | CPU: Itanium2 "Montvale" OS: HP-UX 11iv2,v3, Windows, Linux, OpenVMS | New 4P Server & Chipset CPU: Tukwila OS: HP-UX, Windows Linux, OpenVMS | |
|  HP Integrity rx2620 HP-UX, Windows, Linux, OpenVMS | New 2p/4c Capacity Optimized Server & Chipset CPU: Itanium2 "Montecito" DDR-II PCI-E OS: HP-UX, Windows, Linux, OpenVMS<br>CPU: Itanium2 "Montecito" OS: HP-UX 11iv2, v3*, Windows, Linux, OpenVMS | CPU: Itanium2 "Montvale" OS: HP-UX 11iv2,v3, Windows, Linux, OpenVMS | New Servers & Chipset CPU: Tukwila OS: HP-UX, Windows, Linux, OpenVMS | |
|  HP Integrity rx1620 HP-UX, Windows, Linux, OpenVMS | New 2p/4c Density Optimized Server CPU: Itanium2 "Montecito" DDR-II OS: HP-UX 11iv2, v3, Windows, Linux, OpenVMS PCI-E | CPU: Itanium2 "Montvale" OS: HP-UX 11iv2,v3, Windows, Linux, OpenVMS | | |
|  HP Integrity BL60p Blades | | New Blade Server CPU: LV Montecito OS: HP-UX, Windows*, Linux, OpenVMS* | CPU: Itanium2 "Montvale" OS: HP-UX 11iv2,v3, Windows, Linux, OpenVMS | New Blade Server CPU: Tukwila OS: HP-UX, Windows, Linux. OpenVMS |

Timeframes not to scale

Plans subject to change

All upgrades "in-box" except as noted

*Not available at initial processor release

New Chassis | PCI-Express | DDR-II Memory

# Montecito-based HP Integrity Servers

# Upgrade Paths to Montecito from rx4640, rx2620, and rx1620

## rx4640 to rx4640/Montecito

Processor swap and firmware update

**No Change:** Chassis, system board, chipset, FSB, I/O backplane, disk backplane, power supply, disks, memory and most I/O cards


## rx2620 to rx2620/Montecito

Processor swap and firmware update

**No Change:** Chassis, system board, chipset, FSB, I/O backplane, disk backplane, power supply, disks, memory and most I/O cards can be carried forward

## NO rx1620 to rx1640/Montecito Upgrade





invent

# Supported Montecito SKUs on zx1 systems

- rx2620 supports 2 Montecito processor SKUs starting with a 1.4GHz/12MB dual core (SR mid September).
- The 1.6GHz/18MB SKU will ship release on the rx4640 mid-September. However, availability will be limited at first.

- rx4640 supports 2 Montecito processor SKUs including the 1.6GHz/18MB dual core and the 1.6GHz/24MB dual core. Both 1.6GHz dual core processors will ship release on the rx4640 mid-September.

# Montecito upgrades for zx1 systems

- Current rx4640 and rx2620 customers can easily upgrade their existing zx1 single-core Itanium 2 based systems. Upgrades are customer installable and the dual-core processors are supported on the existing zx1 chipset.
  - rx4640: requires processor swap and firmware upgrade.
  - rx2620: requires processor swap, upgrade kit (baffles & cables) and firmware upgrade

# NEW Entry-class HP Integrity Servers and Chipset in Q3 2006

New in Q3 2006

**New 4-processor Server**

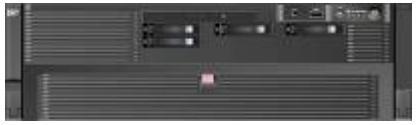

**New 2-processor Server**

** Universal SAS HDDs are identical to those used in HP ProLiant servers and HP StorageWorks products*

## New 4-processor Server

- Up to 4 Intel Itanium® 2 dual-core processors
- New zx2 chipset
- 400 MT/s Front Side Bus
- Up to 192GB of memory
- 2 Gigabit LAN ports and 2 internal SAS channels
- Up to 16 internal, hot plug, universal* hard disk drives
- 8 available PCI-X I/O slots
- 7U rack-optimized or stand-alone form factor
- Operating systems
  - HP-UX 11i v2 and v3 (H2 2006)
  - Red Hat AS 4 & SUSE SLES 10
  - Windows 2003 Enterprise Ed.
  - OpenVMS v8.3

## New 2-processor Server

- Up to 2 Intel Itanium® 2 dual-core processors
- New zx2 chipset
- 533 MT/s Front Side Bus
- Up to 64GB of memory
- 2 Gigabit LAN ports and 2 internal SAS channels
- Up to 8 internal, hot plug, universal* hard disk drives
- 8 available PCI-X I/O slots
- 4U rack-optimized or stand-alone form factor
- Operating systems
  - HP-UX 11i v2 and v3 (H2 2006)
  - Red Hat AS 4 & SUSE SLES 10
  - Windows 2003 Enterprise Ed.
  - OpenVMS v8.3

# rx3600 (Ruby) Features Overview

## Management

- Integrated Lights Out (iLO)
- iLO (integrated Remote Management) Advanced Pack firmware license option

## I/O subsystem

- 8 PCI-X 266MHz slots at SR
  - Upgrade to PCI-E late 2006
- 2 SAS (Serial Attached SCSI) Channels
- 2 x 1 Gigabit ports
- 100Base-T USB, VGA, serial ports

## Processors and chipset

- 1 or 2 Intel ® Montecito dual-core processors
- 1.4GHz/12M or 1.6GHz/18M
- HP zx2 Chipset
- Upgrade to Montvale ~July FY07

## Form factor

- 4 EIA units (U) or 7″ height
- 10 servers per 2m rack
- Designed for data center and utility closet operation (5–35°C)
- Standalone, pedestal option

## High availability

- Redundant hot-plug power
- Redundant hot-plug fans
- Hot-plug PCI-X slots
- Dual SAS channels
- CPU de-allocation on failure
- Dynamic processor resilience

## Internal peripherals

- 8 hot-plug SFF SAS/SATA HDDs
- DVD or DVD/CD-RW

## Memory

- 1 GB to 48 GB (96GB in Q4/06)
- PC2100 ECC chip spare DDR2
- 128-bit, 266 MHz DDR
- 2x12 DIMM card carriers max (does not support 24 DIMM carrier)

## Operating systems

- HP-UX
- Linux Red Hat and SuSE
- Windows Server 2003 Ent. Edition
- OVMS

# rx6600 (Sapphire) Features Overview

## Management

- Integrated Lights Out (iLO)
- iLO (integrated Remote Management) Advanced Pack firmware license option

## I/O subsystem

- 8 PCI-X 266MHz slots at SR
  - Upgrade to PCI-E Q2 FY07
- 2 SAS (Serial Attached SCSI) Channels
- 2 x 1 Gigabit ports
- 100Base-T USB, VGA, serial ports

## Internal peripherals

- 16 hot-plug SFF SAS/SATA HDDs
- DVD or DVD/CD-RW

System Dimensions:

- Height 7U/316.2mm
- Width 440mm
- Depth 672mm

## Processors and chipset

- 1 to 4 Intel ® Montecito dual-core processors
- 1.4GHz/12M, 1.6GHz/18M and 1.6GHz/24M
- HP zx2 Chipset
- Upgrade to Montvale ~July FY07

## Memory

- 1 GB to 96 GB (192GB in Q4/06)
- PC2100 ECC chip spare DDR2
- 128-bit, 266 MHz DDR
- 2x24 DIMM card carriers max (does not support 12 DIMM carrier)

## Form factor

- 7 EIA units (U) or 12.25″ height
- 7 servers per 2m rack
- Designed for data center and utility closet operation (5–35°C)
- Standalone, pedestal option

## High availability

- Redundant hot-plug power
- Redundant hot-plug fans
- Hot-plug PCI-X slots
- Dual SAS channels
- CPU de-allocation on failure
- Dynamic processor resilience

## Operating systems

- HP-UX
- Linux Red Hat and SuSE
- Windows Server 2003 Ent. Edition
- OVMS

# Processor Configuration Rules

- Processors can be installed one at a time
- Processors must be installed in the following sequence: 0, 1, 2, and 3
- Different speed and different cache-size processors cannot be mixed in the same system
- rx3600 will support the 1.4GHz/12M and 1.6GHz/18M at release
- rx6600 will support the 1.4GHz/12M, 1.6GHz/18M and the 1.6GHz/24M

# rx3600 & rx6600 Memory Loading Rules

- Memory must be installed in groups of four DIMMs, also known as quads
- Each quad must consist of equal density DIMMs
- Memory can be ordered in quads of 2 GB (4×512MB), 4 GB (4×1GB), 8 GB (4×2GB), or 16 GB (4×4GB)
- Minimum memory is 2 GB (4×512MB)
- Maximum memory in rx3600 is 96 GB, using six 16-GB memory quads in memory carrier option AD125A (Q4/06)
- Maximum memory in rx6600 is 192 GB, using twelve 16-GB memory quads in memory carrier option AD127A (Q4/06)
- Memory must be loaded in the order depicted on the memory carrier board
- Arrange DIMMs so that the quads with the largest capacity are in the lowest numbered slots.
- For best performance, all DIMM slots on both memory carrier boards should be populated with the same size DIMM.

# rx2660 (Merlion)

01.2007

## Management

- Integrated Lights Out (iLO)
- iLO (integrated Remote Management) Advanced Pack firmware license option

## I/O subsystem

- 3 PCI-Express full-length slots
- 2 SAS (Serial Attached SCSI) Channels
- 4 x 1 Gigabit ports
- 100Base-T USB, VGA, serial ports

## Internal peripherals

- 8 hot-plug SFF SAS/SATA HDDs
- DVD or DVD/CD-RW

## Memory

- 1 GB to 32 GB
- PC2100 ECC chip spare DDR2
- 128-bit, 266 MHz DDR

## Processors and chipset

- 1 or 2 Intel ® Montecito dual-core processors:
- Hi speed bin, Mid-speed bin and
- Lo speed bin
- HP zx2 Chipset

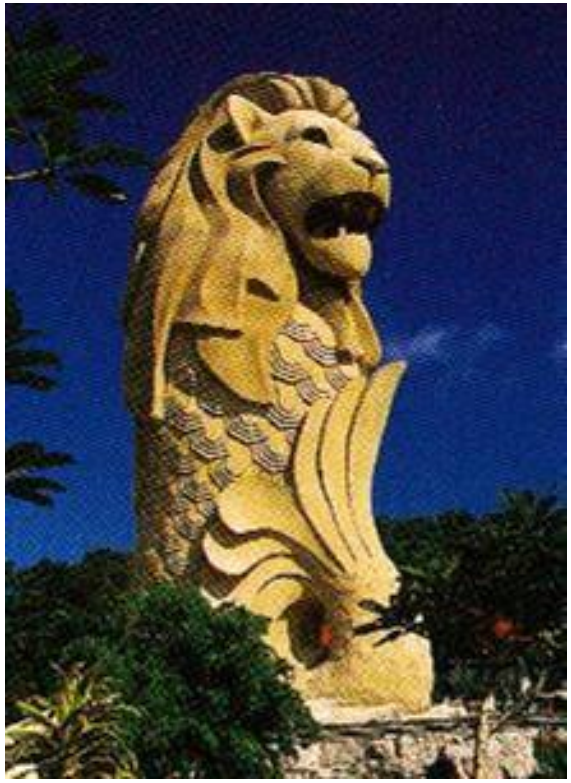

## Form factor

- 2 EIA units (U) or 3.5″ height
- 20 servers per 2m rack
- Designed for data center and utility closet operation (5–35°C)
- Standalone, pedestal option

## High availability

- Redundant hot-plug power
- Redundant hot-plug fans
- Dual SCSI channels
- CPU de-allocation on failure
- Dynamic processor resilience

## Operating systems

- HP-UX
- Linux Red Hat and SuSE
- Windows Server 2003 Ent. Edition
- OVMS

# HP Integrity blade BL60p

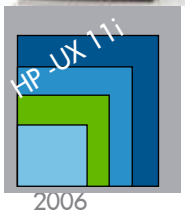


No Montecito Upgrade!!

**Processor**

- One or two Itanium2™ CPUs: 1.6GHz, 3MB L2 cache, 100W power consumption each No Montecito Upgrade
- FSB400 (Mega transfers/sec) 200 MHz Front-Side Bus; 6.4GB/s peak data bandwidth

**Memory**

- 4 memory DIMM slots, 2 DIMM slots per memory cell with chip spare
- 1.2" height DDR DIMMs; 512MB, 1GB, and 2GB densities; max. 8GB
- 266MTransfers/s data, 8.5 GB/s peak data bandwidth

**I/O**

- Four GigE LAN ports standard
- Two 2Gb/s Fiber channel ports standard
- Two USB 2.0 ports via SUV cable

**Manageability**

- Integrity iLO
- Enhanced server manageability, provided by the standard BCS Management Processor (MP)
- LAN telnet console, Web console, Serial port for local console (SU port)
- Temperature monitoring & fan speed control

**Internal Disk Storage Options**

- Two hot-pluggable Ultra320 SCSI disk bays: 36GB, 73GB, and 146GB densities
- Proliant common disk carrier

**OS**

- HP-UX 11i v2, followed by 11i v3



invent

# Next Generation Full Height Blade: BL860c Features Overview

01.2007

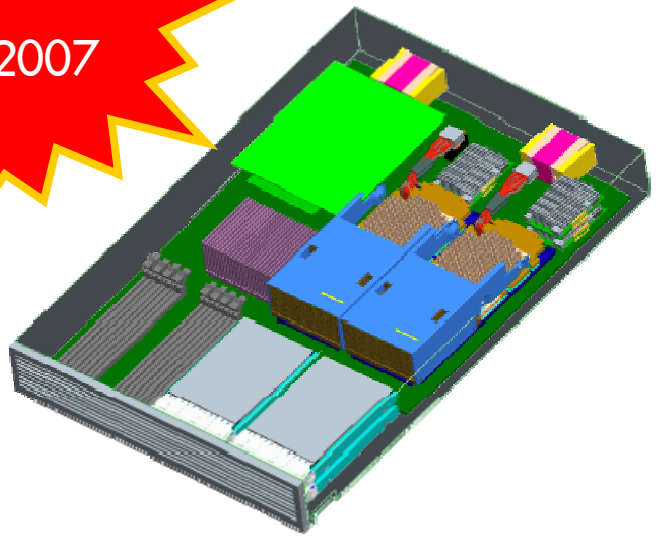

- **Processor**
  - One or two Millington CPUs Dual core, up to 4 CPU cores):
    - 100W power consumption each
  - Perf Price/Perf and Entry SKUs
  - 533 MHz Front-Side Bus 8.6 GB/s peak data bandwidth (upgrade to 667MHz)
- **Memory**
  - 12 memory DIMM slots, 4 DIMM slots per memory cell with double chip kill (48GB max)
  - Standard DDR2 PC2100 DIMMs; 512MB, 1GB, and 2GB and 4GB densities
  - 533MTransfers/s data, 17 GB/s peak data bandwidth
- **System board resident core I/O**
  - SAS Smart Array RAID controller
  - Dual  Dual-port GbE LAN controllers
  - Three standard PCI-E mezzanine expansion slots
  - Two USB 2.0 ports
- **Manageability**
  - Single chip  manageability ASIC and firmware
  - Integrity iLO
  - Overlord standards Front status panel
  - Embedded VGA / 2D graphics display
- **Internal Disk Storage Options**
  - Two  small form factor hot-pluggable SAS disk bays
  - New universal disk carrier

Windows Server2003

HP-UX 11

redhat

simply change

SUSE

- **Operating environments**
  - HP-UX, Linux and Windows Operating systems
  - HP Virtual Machines
  - HP MCSG (Mission Critical ServiceGuard)





hp invent

# sx2000 platforms

## Integrity SuperDome (PA-RISC & IPF)

16 cells, 64 sockets (128 cpu's)
512 DDR-II DIMM slots
192 PCI-X 2.0 slots

New in Q2 2006

## PA-RISC: rp84xx
## IPF: rx86xx

4 cells, 16 sockets (32 cpu's)
64 DDR-II DIMM slots
32 PCI-X 2.0 slots,
17U high, 19" rack

## PA-RISC: rp74xx
## IPF: rx76xx

2 cells, 8 Sockets (16 cpu's)
32 DIMM slots
16 PCI-X 2.0 slots,
10U high, 19" rack

# sx2000 Superdome Family

**32-way**

1 to 4 Cells
2 to 32 Cores
to 512 GB Mem (1 TB cap)
12 to 48 PCI-X slots
1 to 4 nPartitions

**64-way**

1 to 8 Cells
2 to 64 Cores
to 1 TB Mem (2 TB cap)
12 to 48 PCI-X slots (96 w IOX)
1 to 4 nPartitions (8 w IOX)

**128-way**

3 to 16 Cells
6 to 128 Cores
to 2 TB Mem (4TB cap)
36 to 96 PCI-X slots (192 w IOX)
1 to 8 nPartitions (16 w IOX)

# Mixing PA-RISC and Itanium 2 processors within the Integrity Superdome

HP Integrity Superdome

Cell board

supports either

or

sx1000 chipsets

Separate hard partitions support a combination of

Intel Itanium 2 9M
PA-8800
PA-8900

sx2000 chipset

Separate hard partitions support a combination of

Intel Itanium 2 9M
PA-8900
"Montecito"

Processors

# Montecito Processor Loading rules

| Configuration | CPU0 Location | CPU1 Location | CPU3 Location | CPU2 Location |
|---|---|---|---|---|
| 1 Processor | CPU Module | Unpopulated | Unpopulated | Unpopulated |
| 2 Processors | CPU Module | Unpopulated | Unpopulated | CPU Module |
| 3 Processors | CPU Module | CPU Module | Unpopulated | CPU Module |
| 4 Processors | CPU Module | CPU Module | CPU Module | CPU Module |

- Cell Board **does not require** that a **termination module** be installed at the end of an unused FSB. The Arches chipset allows system firmware to disable an unused FSB in the Cache Coherency Controller.

- All Processors installed must be the same type and speed, both on the individual cell as well as in the nPartition. Mixing only allowed in separate nPars.

# SD Memory configuration

- Memory module sizes are now comprised of 8 DIMMs
  - 8 GB module = eight 1 GB DIMMs
  - 16 GB module = eight 2 GB DIMMs
  - 32 GB module = eight 4 GB DIMMs (coming with Montecito)
- Memory Modules of different size can be mixed on a single cell board
- Larger DIMMs must be loaded first across a cell, followed by progressively smaller DIMM sizes
- For best performance populate each cell in an nPar with 4 Memory modules of equal size.

# Three most important rules to follow

1. Balance your memory
   - Load DIMMs evenly across all four memory subsystems
     - –Populate all SDRAM buses if possible (requires 16 DIMMs)
     - –More DIMMs is better for performance
       - –sixteen - 1GB DIMMs beats eight – 2GB DIMMs every time
   - Load memory evenly across all the Cells in a hard partition
     - –Each Cell in the nPar should have same amount of memory
2. Use power of 2 number of Cells (1,2,4,8,16) for partition sizes whenever possible
   - – Global Interleaving assures even distribution of memory accesses
   - – Minimal "inaccessible" memory loss (only 64 MB per Cell)
3. Populate memory in power of 2 volumes (8 GB, 16 GB, 32 GB, 64 GB, 128 GB, 256 GB, 512 GB, 1 TB, 2 TB) whenever possible
   - – Maintaining this rule down to the Cell level results in best possible local Cell interleaving characteristics as well

# sx2000 in rx8640

- Two to 32 Itanium2® cores
- Same cellular technology as Superdome
  - One to four cells
  - 64 DIMM slots → 256 GB at first release
    - 512 GB when DRAMs available
  - "Double chipkill" multi-bit error correction
  - 16 PCI-X 266 I/O slots and 2 core I/Os
  - Performance, partitioning, etc.
- System Package
  - 17U, 28.5 inches deep
  - 4 hot-swap hard disks bays (NEW U320 SCSI)
  - 2 removable media bays
  - N+1 OLR fans and bulk power supplies
  - Optional redundant power cords
- Optional I/O expander
  - 16 additional PCI-X I/O slots
  - 4 more hot-swap disk and 2 more removable media bays

# sx2000 in rx8640

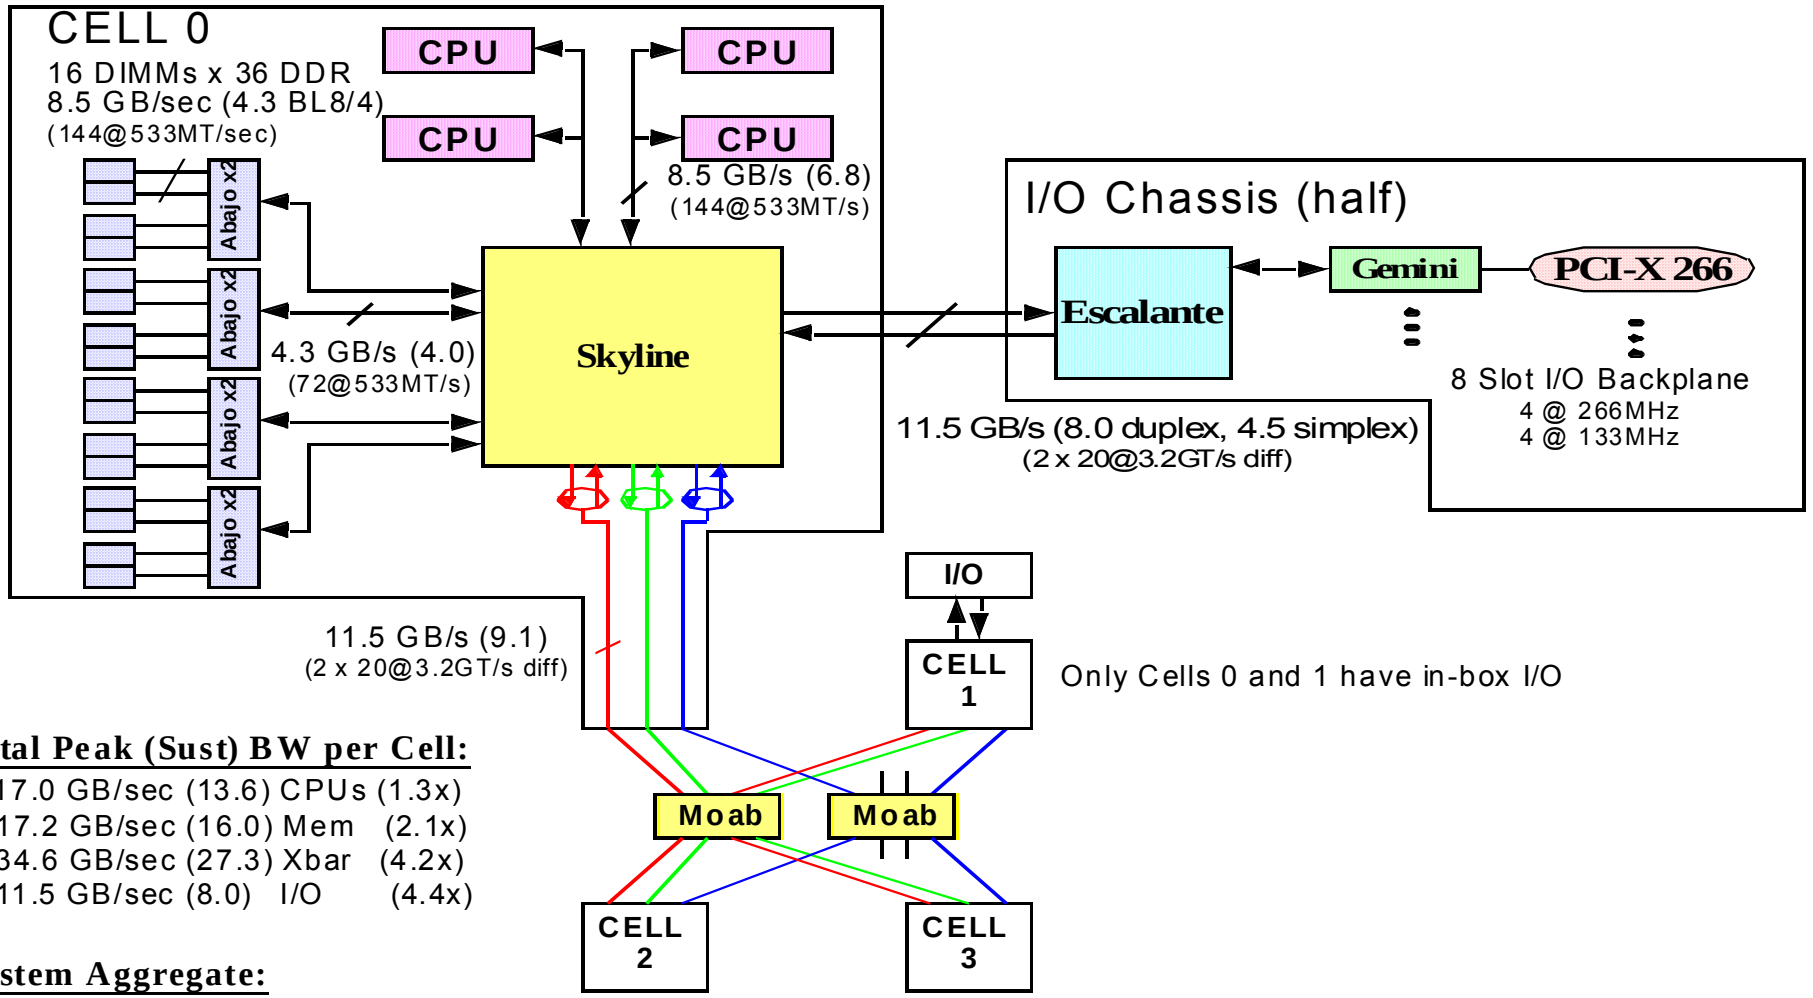


**Total Peak (Sust) BW per Cell:**

17.0 GB/sec (13.6) CPUs (1.3x)
17.2 GB/sec (16.0) Mem (2.1x)
34.6 GB/sec (27.3) Xbar (4.2x)
11.5 GB/sec (8.0) I/O (4.4x)

**System Aggregate:**

>250 GFlops/sec with Montecito
55 GB/sec bisection BW
32 GB/sec I/O BW: >2,750 TB/day



invent

# sx2000 in rx7640

- Two to 16 Itanium2® cores
- Same cellular technology as Superdome
  - One to two cells
  - 32 DIMM slots → 128 GB at first release
    - 256 GB when DRAMs available
  - "Double chipkill" multi-bit error correction
  - 16 PCI-X 266 I/O slots
  - Performance, partitioning, etc. of Superdome
- System Package
  - 10U, 28.5 inches deep
  - 4 hot-swap hard disks bays (NEW each disk now on separate SCSI bus)
  - 1 removable media bay with 2 slim-line DVD
  - N+1 OLR fans, 2N bulk power supplies
  - Optional redundant power cords

# sx2000 in rx7640

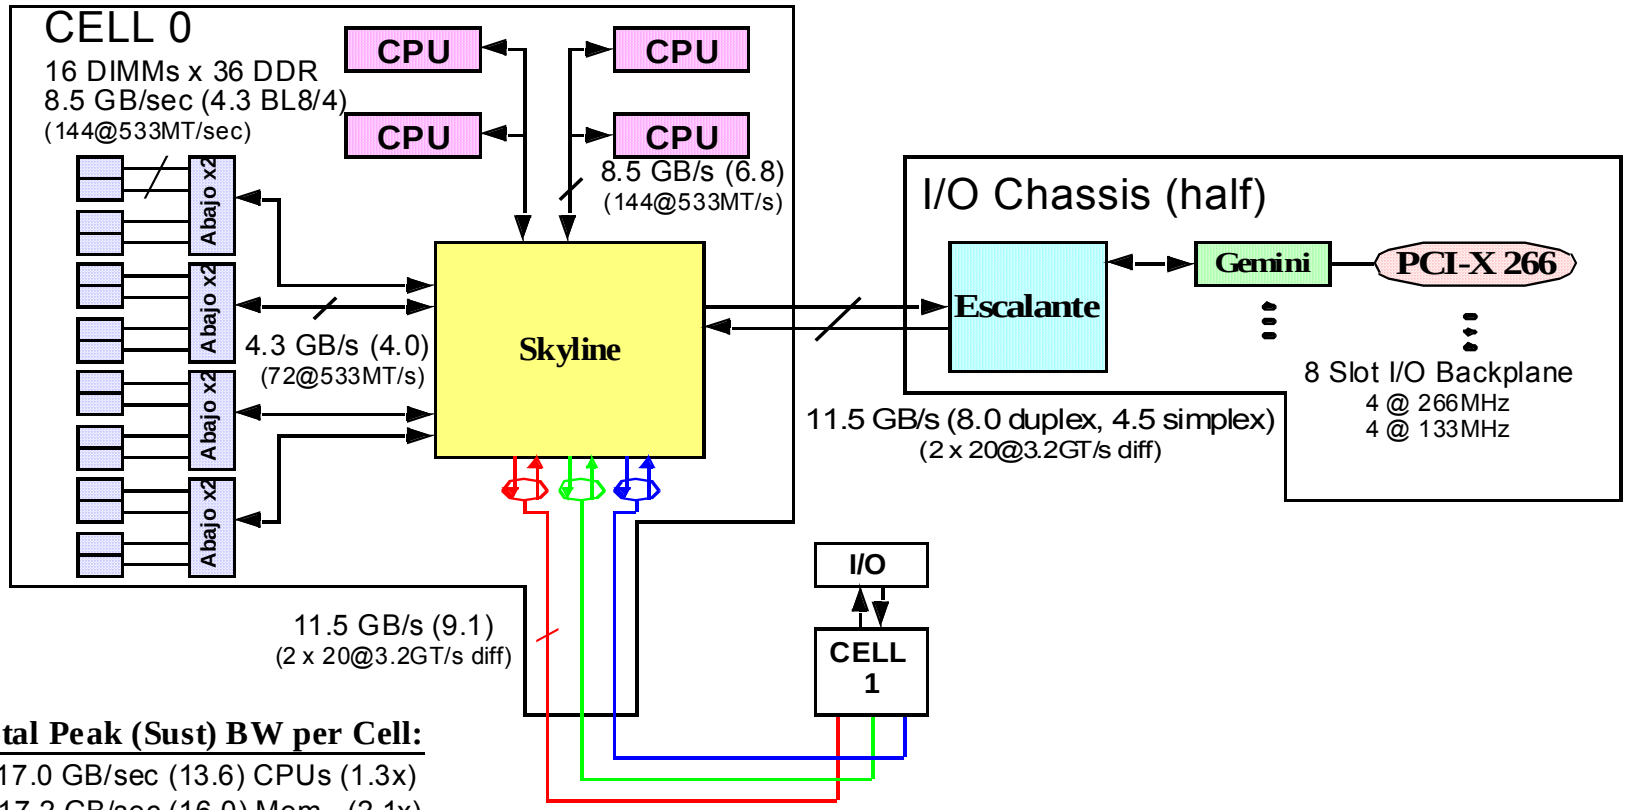


**Total Peak (Sust) BW per Cell:**

17.0 GB/sec (13.6) CPUs (1.3x)
17.2 GB/sec (16.0) Mem (2.1x)
34.6 GB/sec (27.3) Xbar (4.2x)
11.5 GB/sec (8.0) I/O (4.4x)

**System Aggregate:**

>125 GFlops/sec with Montecito
27 GB/sec bisection BW
16 GB/sec I/O BW: >1,375 TB/day

# Montecito Processor Loading rules

| Configuration | CPU0 Location | CPU1 Location | CPU3 Location | CPU2 Location |
|---|---|---|---|---|
| 1 Processor | CPU Module | Unpopulated | Unpopulated | Unpopulated |
| 2 Processors | CPU Module | Unpopulated | Unpopulated | CPU Module |
| 3 Processors | CPU Module | CPU Module | Unpopulated | CPU Module |
| 4 Processors | CPU Module | CPU Module | CPU Module | CPU Module |

- Cell Board **does not require** that a **termination module** be installed at the end of an unused FSB. The Arches chipset allows system firmware to disable an unused FSB in the Cache Coherency Controller.

- All Processors installed must be the same type and speed, both on the individual cell as well as in the nPartition. Mixing only allowed in separate nPars.

# Mid-Range memory

- In the memory area, until September 1, 2006, the rx7640 and rx8640 sx2000 based chassises will support only ½ of the rx7620/rx8620 maximum memory…64-GBs (rx7640)/128-GBs (rx8640).
- Memory sizes are a 2-GB (two 1-GB DIMMs) and a 4-GB module (two 2-GB DIMMs) added in July 2006.
- An 8-GB module (two 4-GB DIMMs) will be supported on September 1, 2006.
- Please note that memory modules in the rx7640 and rx8640 are two DIMMs, not four as in the rx7620 and rx8620 servers.  Memory modules can be mixed on a cell board.

# rx8640 and rx7640 Memory Configurations

Memory module sizes are now comprised of 2 DIMMs
- 2 GB module = two 1 GB DIMMs
- 4 GB module = two 2 GB DIMMs
- 8 GB module = two 4 GB DIMMs (coming with Montecito)

- Since the combinations aren't limited by the module size (as with Superdome), these mid-range platforms have additional configuration rules

- Why are there two different module sizes for the high-end and mid-range products ?
  - Different products – different market segments – different needs
  - Entry pricing pressure on mid-range forced module size down
  - High-end, even with larger module size, still offers more memory sizing flexibility than our competition, and configurations which do not perform well are eliminated

# rx8640 and rx7640 Memory Configurations (continued)

- rx8640 and rx7640 Memory configuration Rules
  - DIMMs must be loaded in pairs (same size within a pair)
  - DIMM pairs must be loaded in slot order
  - Larger DIMMs must be loaded first across a cell, followed by progressively smaller DIMM sizes
- The table below gives a few examples of supported configurations, but is not comprehensive

| Memory per Cell (GBytes) | Number of DIMMs | | | Bandwidth (estimate) | Echelon | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | 1 GB | 2 GB | 4 GB | | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| 2 | 2 | | | very poor | | | | | | | | |
| 4 | 4 | | | quite poor | | | | | | | | |
| 8 | 8 | | | average | | | | | | | | |
| **16** | **16** | | | **best** | | | | | | | | |
| 24 | 8 | 8 | | good | | | | | | | | |
| **32** | | **16** | | **best** | | | | | | | | |
| 48 | | 8 | 8 | good | | | | | | | | |
| **64** | | | **16** | **best** | | | | | | | | |

# Supported OS on mid range and high end server with Montecito

| Operating System | Revision | Orderable with Field Installation | Orderable with Fact. Integration (opt. 0D1) | Shippable with Field Installation | Shippable with Fact. Integration (opt. 0D1) |
|---|---|---|---|---|---|
| HP-UX | 11I V2<br>Fusion 0609 | August 1 | August 1 | at SR | at SR. |
| vPar | A.04.02 | | | | |
| VM | Version 1.2 | | | | |
| Windows 2003 | SP1 | August 1 | August 1 | at SR | at SR |
| Linux RedHat | RHEL4 U3 | No* | TBD*<br>(RHEL4U4) | No* | TBD*<br>(RHEL4U4) |
| Linux SUSE | SELS 10 | rx8640 Sep.1<br>rx7640 Dec.1<br><br>SD Dec.1 | Mid-Range<br>Not offered<br><br>SD Dec.1 | rx8640 at SR<br>rx7640 Dec.<br><br>SD Dec. | Mid-Range<br>Not offered<br><br>SD Dec. |
| OpenVMS | V8.3 | August 1 | August 1 | at SR | at SR |

* RHEL4 U3 will not work on all Montecito based servers (vendor independent). It is being worked on with RedHat to have RHEL4 U4 supported ASAP. This is not an issue with the SUSE release.

# Multi-Core Licensing

- HP-UX and layered SW = 1 license/core
- Windows 2003 Server & SQL 2005 = 1 license/processor
- Oracle
  - multi-core RISC processors: 0.75 licenses/core
    - PA-RISC, IBM Power5, SUN UltraSparcIV
  - Intel and AMD multi-core processors: 0.5 licenses/core
    - x86 and Itanium
  - SUN T1 : 0.25 licenses/core
- BEA = 1 license/processor

# Investment protection with the sx2000 chipset

Integrity Superdome

sx1000 chipset + Itanium 2 9M processors → sx2000 chipset + Itanium 2 9M processors → sx2000 chipset + Itanium 2 "Montecito" processors

rx8620

rx8620-32: sx1000 chipset + Itanium 2 9M processors → rx8640: sx2000 chipset + Itanium 2 9M processors → rx8640: sx2000 chipset + Itanium 2 "Montecito" processors

rx7620

rx7620-32: sx1000 chipset + Itanium 2 9M processors → rx7640: sx2000 chipset + Itanium 2 9M processors → rx7640: sx2000 chipset + Itanium 2 "Montecito" processors

# Upgrading sx2000 Integrity Superdome

- What gets replaced ?
  - Cells (including all memory) and Cell power boards
    - Mad9M CPU modules may be re-usable (TBD)
  - Backplanes and Backplane power boards
    - New Frame design, clock subsystem plugs in Left Backplane
  - I/O Chassis, including all I/O Link cables (many I/O cards still supported)
- Old clock cables are discarded
  - New clock cable design required only for dual-cabinet config
- Flex cable interconnect eliminated
  - New M-Link cables connecting backplanes are more robust, but interconnect is more complicated
- Keep the chassis, fans, bulk power supplies, and the Manageability (Utility) chassis and Utility cables
  - Management Processor must be version 2

# Learning Check

- Q1.What is the name of the cache self healing technology on Montecito?

- A1: Pellston
- A2: Silverdale
- A3: Montvale

# Learning Check

- Q2.Which type of memory is supported with the new sx2000 chipset?

- A1: SDRAM
- A2: RAMBUS-DRAM
- A3: DDR-2 memory


invent